



US009195517B2

(12) **United States Patent**
Faraj

(10) **Patent No.:** **US 9,195,517 B2**
(45) **Date of Patent:** ***Nov. 24, 2015**

(54) **DETERMINING COLLECTIVE BARRIER
OPERATION SKEW IN A PARALLEL
COMPUTER**

(56) **References Cited**
U.S. PATENT DOCUMENTS

(71) Applicant: **INTERNATIONAL BUSINESS
MACHINES CORPORATION,**
Armonk, NY (US)

5,499,349	A	3/1996	Nikhil et al.
5,535,387	A	7/1996	Matsuoka et al.
5,740,067	A	4/1998	Hathaway
6,117,180	A	9/2000	Dave et al.
7,958,285	B1	6/2011	Chiu et al.
8,234,652	B2	7/2012	Arimilli et al.

(72) Inventor: **Daniel A. Faraj,** Rochester, MN (US)

(73) Assignee: **International Business Machines
Corporation,** Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 342 days.

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **13/685,869**

(22) Filed: **Nov. 27, 2012**

(65) **Prior Publication Data**

US 2013/0145379 A1 Jun. 6, 2013

Related U.S. Application Data

(63) Continuation of application No. 13/308,917, filed on
Dec. 1, 2011.

(51) **Int. Cl.**
G06F 9/46 (2006.01)
G06F 15/16 (2006.01)
G06F 9/52 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 9/522** (2013.01)

(58) **Field of Classification Search**
None
See application file for complete search history.

(Continued)
OTHER PUBLICATIONS

Jones et al., "A Clock Synchronization Strategy for Minimizing
Clock Variance At Runtime in High-End Computing Environments",
2010 22nd International Symposium on Computer Architecture and
High Performance Computing (SBAC-PAD), Oct. 27-30, 2010, pp.
207-214, Institute of Electrical and Electronics Engineers (IEEE),
Digital Object Identifier: 10.1109/SBAC-PAD.2010.33, USA.

(Continued)

Primary Examiner — Tammy Lee

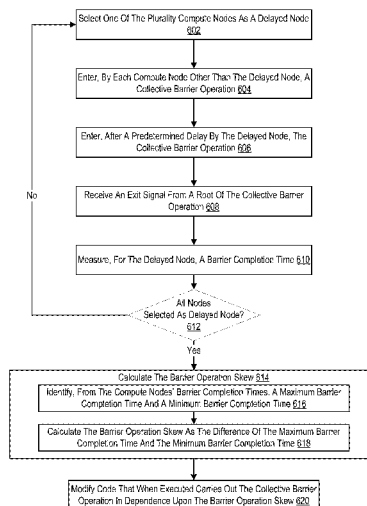
Assistant Examiner — Michael Ayers

(74) *Attorney, Agent, or Firm* — Edward J. Lenart; Grant A.
Johnson; Kennedy Lenart Spraggins LLP

(57) **ABSTRACT**

Determining collective barrier operation skew in a parallel
computer that includes a number of compute nodes organized
into an operational group includes: for each of the nodes until
each node has been selected as a delayed node: selecting one
of the nodes as a delayed node; entering, by each node other
than the delayed node, a collective barrier operation; entering,
after a delay by the delayed node, the collective barrier opera-
tion; receiving an exit signal from a root of the collective barrier
operation; and measuring, for the delayed node, a barrier completion
time. The barrier operation skew is calculated by: identifying, from
the compute nodes' barrier completion times, a maximum barrier
completion time and a minimum barrier completion time and
calculating the barrier operation skew as the difference of the
maximum and the minimum barrier completion time.

4 Claims, 8 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2003/0093569	A1 *	5/2003	Sivier et al.	709/248
2006/0166689	A1	7/2006	Phillips	
2008/0077921	A1	3/2008	Chaudhary et al.	
2008/0270800	A1	10/2008	Krober et al.	
2009/0037707	A1	2/2009	Blocksome	
2009/0046585	A1 *	2/2009	Faraj	370/236
2009/0049323	A1	2/2009	Imark et al.	
2009/0063880	A1	3/2009	Arimilli et al.	
2009/0064140	A1	3/2009	Arimilli et al.	
2009/0064176	A1	3/2009	Ohly et al.	
2009/0248900	A1	10/2009	Maruchek et al.	
2011/0119038	A1	5/2011	Oishi et al.	
2011/0119468	A1	5/2011	Kumar et al.	
2011/0252264	A1	10/2011	Solinas et al.	
2013/0159758	A1	6/2013	Chen et al.	
2013/0159760	A1	6/2013	Chen et al.	

OTHER PUBLICATIONS

Becker et al., "Replay-Based Synchronization of Timestamps in Event Traces of Massively Parallel Applications", International Conference on Parallel Processing—Workshops 2008 (ICPP-W'08),

Sep. 8-12, 2008, pp. 212-219, Institute of Electrical and Electronics Engineers (IEEE), Digital Object Identifier: 10.1109/ICPP-W.2008.17, USA.

Doleschal et al., "Internal Timer Synchronization for Parallel Event Tracing", 2008 Proceedings of the 15th European PVM/MPI Users' Group Meeting on Recent Advances in Parallel Virtual Machine and Message Passing Interface, Sep. 2008, pp. 202-209, vol. 5205, Springer Berlin Heidelberg, NY, Digital Object Identifier: 10.1007/978-3-540-87475-1_29.

Cristian et al., "Probabilistic Internal Clock Synchronization", Proceedings of the Thirteenth Symposium on Reliable Distributed Systems, May 2003, pp. 22-31, California.

Gusella et al., "The Accuracy of the Clock Synchronization Achieved by TEMPO in Berkeley UNIX 4.3BSD", IEEE Transactions on Software Engineering, Jul. 1989, vol. 15, No. 7, pp. 847-853, Institute of Electrical and Electronics Engineers (IEEE), Digital Object Identifier: 10.1109/32.29484, USA.

Office Action, U.S. Appl. No. 13/327,107, Jan. 27, 2014, pp. 1-18. Final Office Action, U.S. Appl. No. 13/693,894, Jun. 26, 2014, pp. 1-16.

Office Action, U.S. Appl. No. 13/693,894, Feb. 5, 2014, pp. 1-15.

* cited by examiner

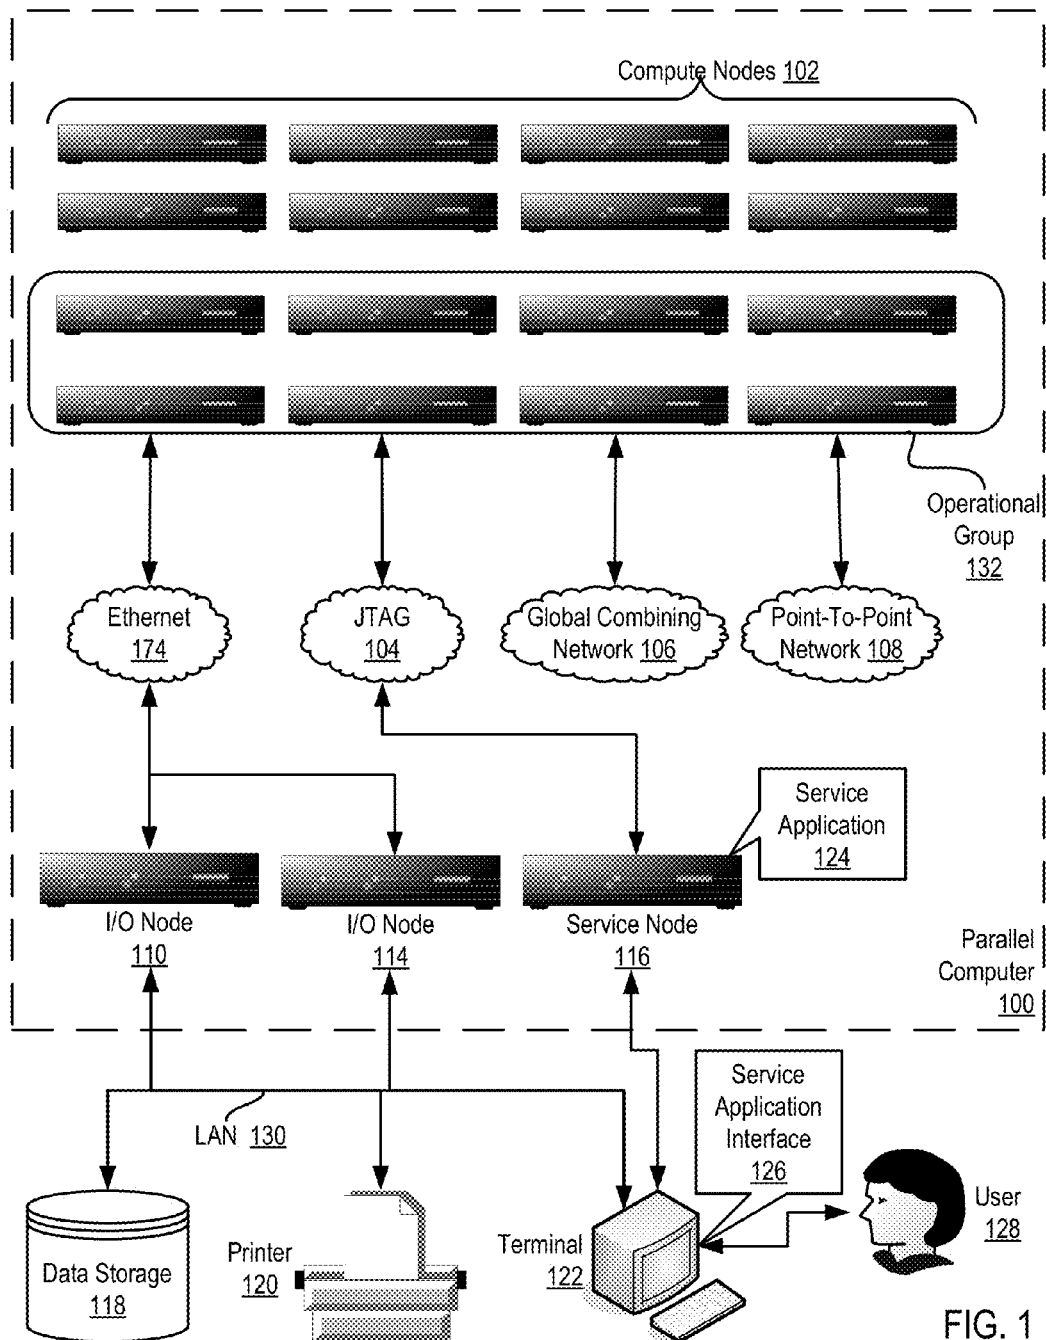


FIG. 1




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4

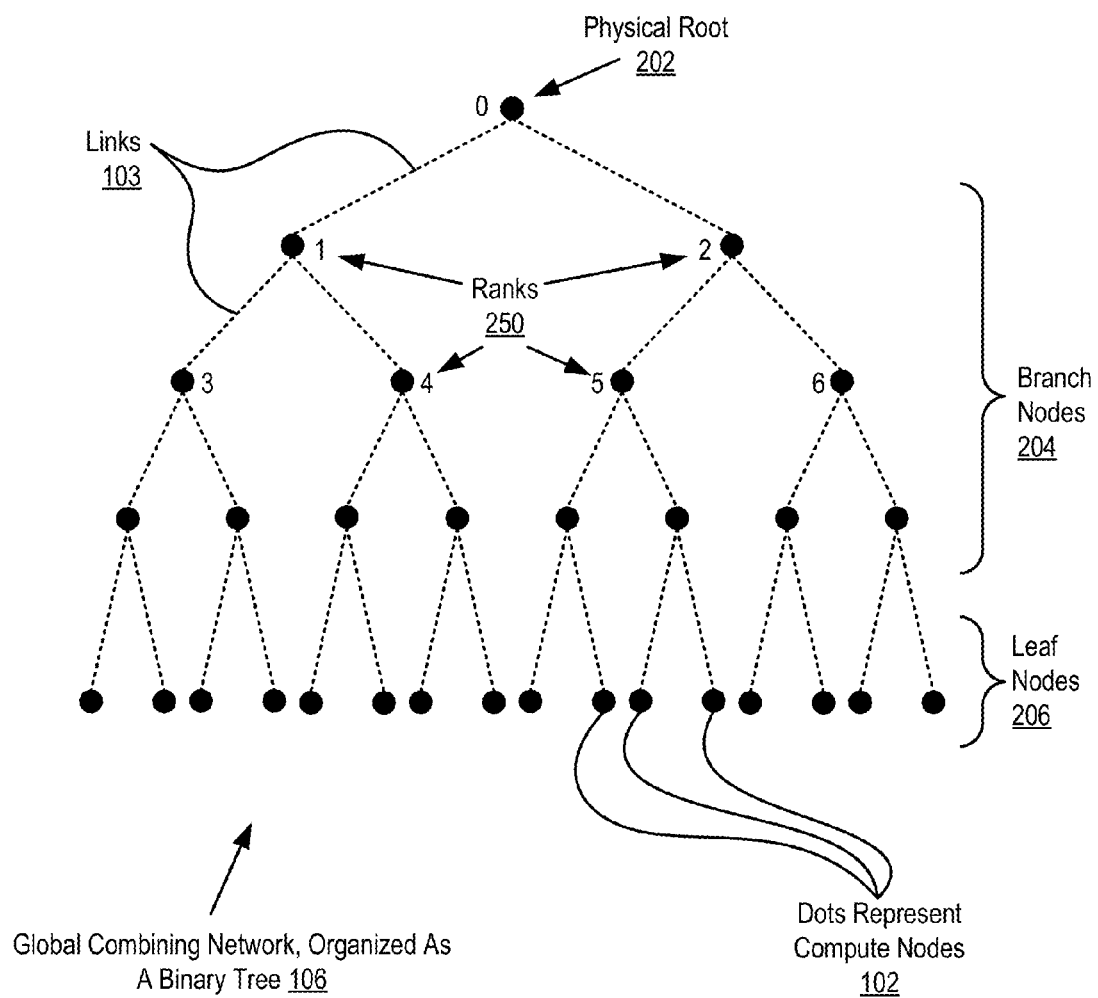


FIG. 5

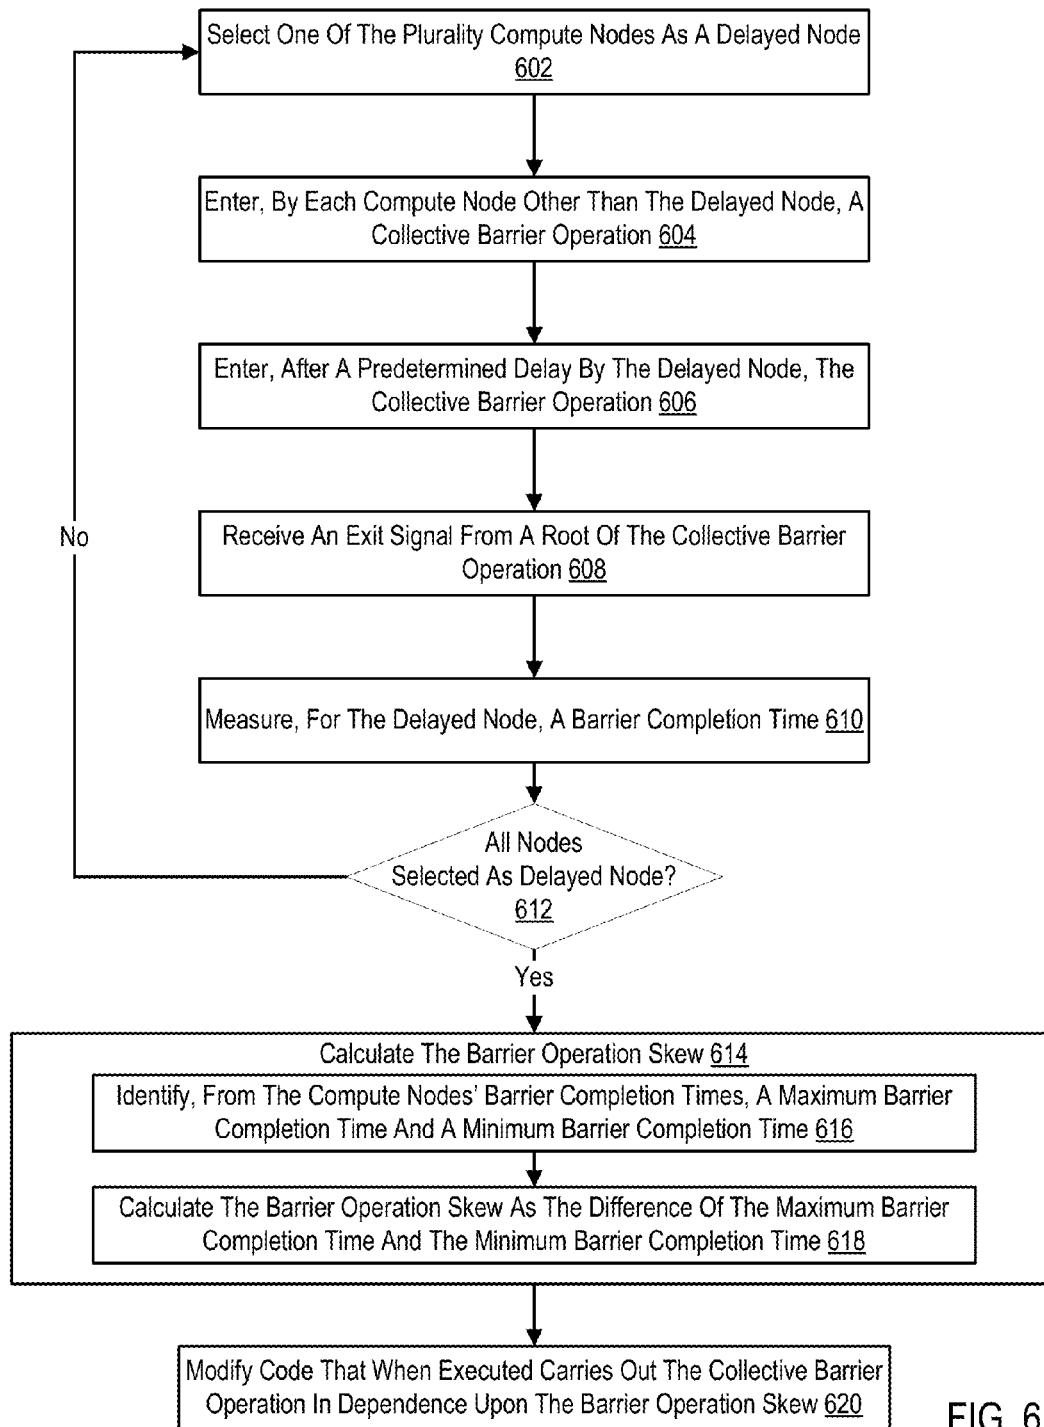


FIG. 6




FIG. 7

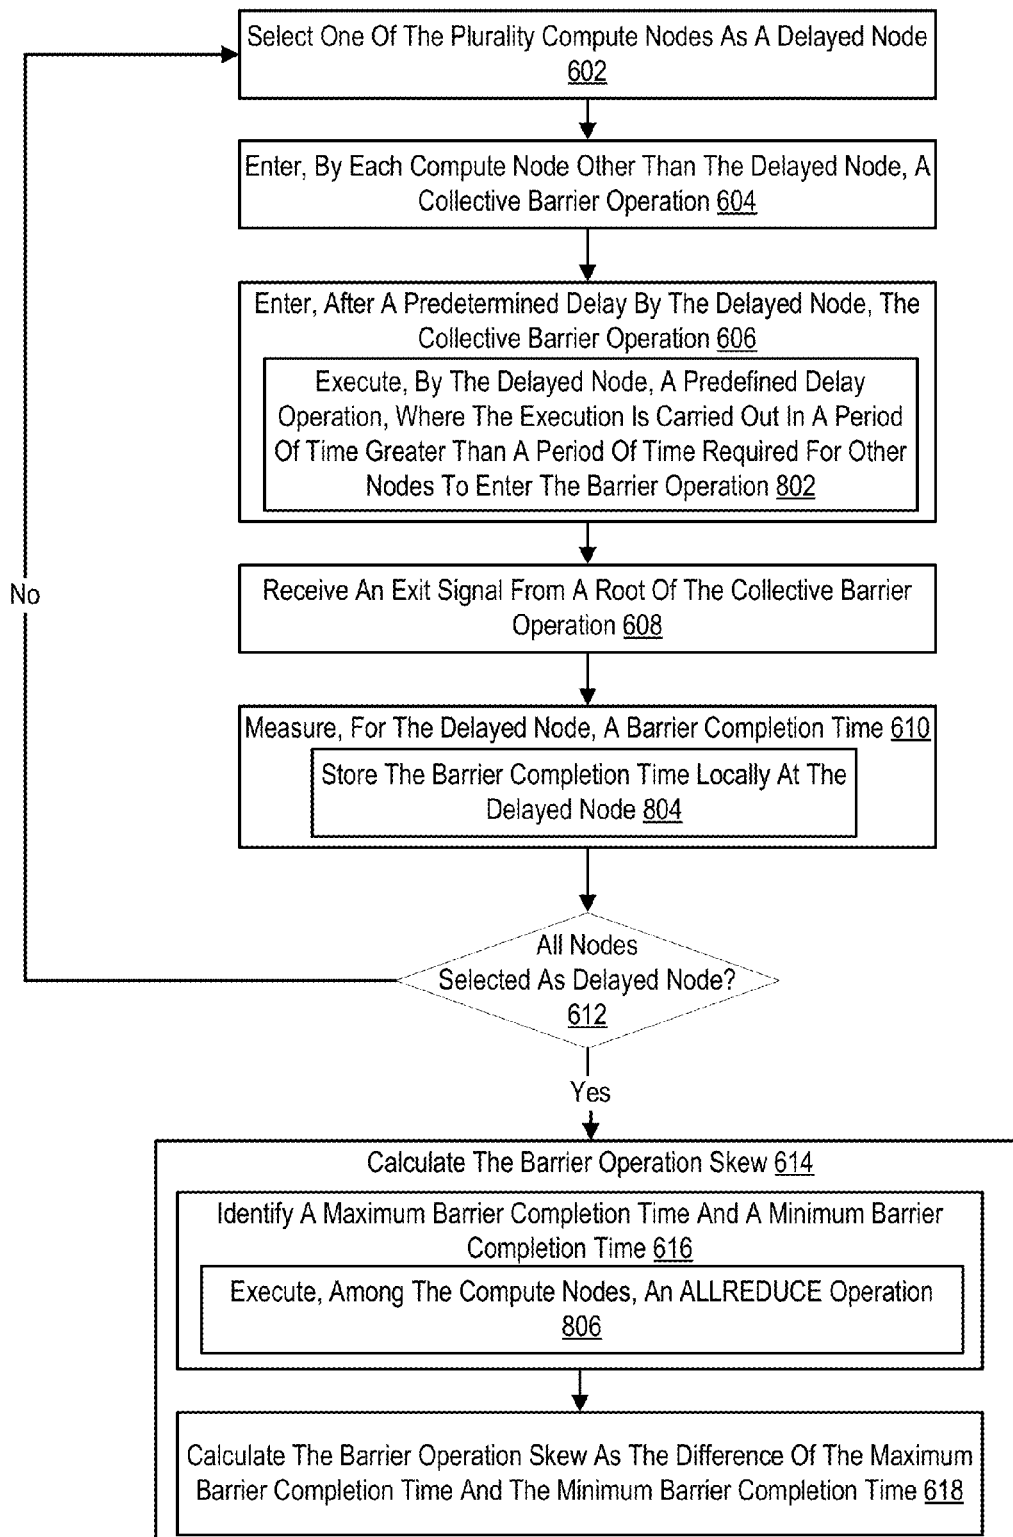


FIG. 8

1

DETERMINING COLLECTIVE BARRIER OPERATION SKEW IN A PARALLEL COMPUTER

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation application of and claims priority from U.S. patent application Ser. No. 13/308,917, filed on Dec. 1, 2011.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

This invention was made with Government support under Contract No. B554331 awarded by the Department of Energy. The Government has certain rights in this invention.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The field of the invention is data processing, or, more specifically, methods, apparatus, and products for determining collective barrier operation skew in a parallel computer.

2. Description of Related Art

The development of the EDVAC computer system of 1948 is often cited as the beginning of the computer era. Since that time, computer systems have evolved into extremely complicated devices. Today's computers are much more sophisticated than early systems such as the EDVAC. Computer systems typically include a combination of hardware and software components, application programs, operating systems, processors, buses, memory, input/output devices, and so on. As advances in semiconductor processing and computer architecture push the performance of the computer higher and higher, more sophisticated computer software has evolved to take advantage of the higher performance of the hardware, resulting in computer systems today that are much more powerful than just a few years ago.

Parallel computing is an area of computer technology that has experienced advances. Parallel computing is the simultaneous execution of the same task (split up and specially adapted) on multiple processors in order to obtain results faster. Parallel computing is based on the fact that the process of solving a problem usually can be divided into smaller tasks, which may be carried out simultaneously with some coordination.

Parallel computers execute parallel algorithms. A parallel algorithm can be split up to be executed a piece at a time on many different processing devices, and then put back together again at the end to get a data processing result. Some algorithms are easy to divide up into pieces. Splitting up the job of checking all of the numbers from one to a hundred thousand to see which are primes could be done, for example, by assigning a subset of the numbers to each available processor, and then putting the list of positive results back together. In this specification, the multiple processing devices that execute the individual pieces of a parallel program are referred to as 'compute nodes.' A parallel computer is composed of compute nodes and other processing nodes as well, including, for example, input/output ('I/O') nodes, and service nodes. Parallel algorithms are valuable because it is faster to perform some kinds of large computing tasks via a parallel algorithm than it is via a serial (non-parallel) algorithm, because of the way modern processors work. It is far more difficult to construct a computer with a single fast processor than one with many slow processors with the same

2

throughput. There are also certain theoretical limits to the potential speed of serial processors. On the other hand, every parallel algorithm has a serial part and so parallel algorithms have a saturation point. After that point adding more processors does not yield any more throughput but only increases the overhead and cost.

Parallel algorithms are designed also to optimize one more resource the data communications requirements among the nodes of a parallel computer. There are two ways parallel processors communicate, shared memory or message passing. Shared memory processing needs additional locking for the data and imposes the overhead of additional processor and bus cycles and also serializes some portion of the algorithm.

Message passing processing uses high-speed data communications networks and message buffers, but this communication adds transfer overhead on the data communications networks as well as additional memory need for message buffers and latency in the data communications among nodes. Designs of parallel computers use specially designed data communications links so that the communication overhead will be small but it is the parallel algorithm that decides the volume of the traffic.

Many data communications network architectures are used for message passing among nodes in parallel computers. Compute nodes may be organized in a network as a 'torus' or 'mesh,' for example. Also, compute nodes may be organized in a network as a tree. A torus network connects the nodes in a three-dimensional mesh with wrap around links. Every node is connected to its six neighbors through this torus network, and each node is addressed by its x,y,z coordinate in the mesh. In such a manner, a torus network lends itself to point to point operations. In a tree network, the nodes typically are connected into a binary tree: each node has a parent and two children (although some nodes may only have zero children or one child, depending on the hardware configuration). Although a tree network typically is inefficient in point to point communication, a tree network does provide high bandwidth and low latency for certain collective operations, message passing operations where all compute nodes participate simultaneously, such as, for example, an allgather operation. In computers that use a torus and a tree network, the two networks typically are implemented independently of one another, with separate routing circuits, separate physical links, and separate message buffers.

SUMMARY OF THE INVENTION

Methods, apparatus, and products for determining collective barrier operation skew in a parallel computer are disclosed in this specification. The parallel computer includes a plurality of compute nodes organized into an operational group and determining collective barrier operation skew includes an iterative process carried out once for each of the compute nodes until each compute node has been selected as a delayed node. The iterative process includes: selecting one of the plurality compute nodes as a delayed node; entering, by each compute node other than the delayed node, a collective barrier operation; entering, after a predetermined delay by the delayed node, the collective barrier operation; receiving an exit signal from a root of the collective barrier operation; and measuring, for the delayed node, a barrier completion time, where the barrier completion time is the period of time between the delayed node entering the barrier and receiving the exit signal. Once each compute node in the operational group has measured a barrier completion time, the barrier operation skew may be calculated by: identifying, from the compute nodes' barrier completion times, a maximum barrier

completion time and a minimum barrier completion time and calculating the barrier operation skew as the difference of the maximum barrier completion time and the minimum barrier completion time.

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular descriptions of exemplary embodiments of the invention as illustrated in the accompanying drawings wherein like reference numbers generally represent like parts of exemplary embodiments of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an exemplary system for determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 2 sets forth a block diagram of an example compute node useful in a parallel computer capable of determining collective barrier operation skew according to embodiments of the present invention.

FIG. 3A sets forth a block diagram of an example Point-To-Point Adapter useful in systems for determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 3B sets forth a block diagram of an example Global Combining Network Adapter useful in systems for determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 4 sets forth a line drawing illustrating an example data communications network optimized for point-to-point operations useful in systems capable of determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 5 sets forth a line drawing illustrating an example global combining network useful in systems capable of determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 6 sets forth a flow chart illustrating an example method for determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 7 sets forth a line drawing illustrating an example operational group of compute nodes useful in determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

FIG. 8 sets forth a flow chart illustrating a further example method determining collective barrier operation skew in a parallel computer according to embodiments of the present invention.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

Exemplary methods, apparatus, and products for determining collective barrier operation skew in a parallel computer in accordance with the present invention are described with reference to the accompanying drawings, beginning with FIG. 1. FIG. 1 illustrates an exemplary system for determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. The system of FIG. 1 includes a parallel computer (100), non-volatile memory for the computer in the form of a data storage device (118), an output device for the computer in the form of a printer (120), and an input/output device for the computer in the form of a computer terminal (122).

The parallel computer (100) in the example of FIG. 1 includes a plurality of compute nodes (102). The compute

nodes (102) are coupled for data communications by several independent data communications networks including a high speed Ethernet network (174), a Joint Test Action Group ('JTAG') network (104), a global combining network (106) which is optimized for collective operations using a binary tree network topology, and a point-to-point network (108), which is optimized for point-to-point operations using a torus network topology. The global combining network (106) is a data communications network that includes data communications links connected to the compute nodes (102) so as to organize the compute nodes (102) as a binary tree. Each data communications network is implemented with data communications links among the compute nodes (102). The data communications links provide data communications for parallel operations among the compute nodes (102) of the parallel computer (100).

The compute nodes (102) of the parallel computer (100) are organized into at least one operational group (132) of compute nodes for collective parallel operations on the parallel computer (100). Each operational group (132) of compute nodes is the set of compute nodes upon which a collective parallel operation executes. Each compute node in the operational group (132) is assigned a unique rank that identifies the particular compute node in the operational group (132). Collective operations are implemented with data communications among the compute nodes of an operational group. Collective operations are those functions that involve all the compute nodes of an operational group (132). A collective operation is an operation, a message-passing computer program instruction that is executed simultaneously, that is, at approximately the same time, by all the compute nodes in an operational group (132) of compute nodes. Such an operational group (132) may include all the compute nodes (102) in a parallel computer (100) or a subset all the compute nodes (102). Collective operations are often built around point-to-point operations. A collective operation requires that all processes on all compute nodes within an operational group (132) call the same collective operation with matching arguments. A 'broadcast' is an example of a collective operation for moving data among compute nodes of an operational group. A 'reduce' operation is an example of a collective operation that executes arithmetic or logical functions on data distributed among the compute nodes of an operational group (132). An operational group (132) may be implemented as, for example, an MPI 'communicator.'

'MPI' refers to 'Message Passing Interface,' a prior art parallel communications library, a module of computer program instructions for data communications on parallel computers. Examples of prior-art parallel communications libraries that may be improved for use in systems configured according to embodiments of the present invention include MPI and the 'Parallel Virtual Machine' ('PVM') library. PVM was developed by the University of Tennessee, The Oak Ridge National Laboratory and Emory University. MPI is promulgated by the MPI Forum, an open group with representatives from many organizations that define and maintain the MPI standard. MPI at the time of this writing is a de facto standard for communication among compute nodes running a parallel program on a distributed memory parallel computer. This specification sometimes uses MPI terminology for ease of explanation, although the use of MPI as such is not a requirement or limitation of the present invention.

Some collective operations have a single originating or receiving process running on a particular compute node in an operational group (132). For example, in a 'broadcast' collective operation, the process on the compute node that distributes the data to all the other compute nodes is an originat-

ing process. In a 'gather' operation, for example, the process on the compute node that received all the data from the other compute nodes is a receiving process. The compute node on which such an originating or receiving process runs is referred to as a logical root.

Most collective operations are variations or combinations of four basic operations: broadcast, gather, scatter, and reduce. The interfaces for these collective operations are defined in the MPI standards promulgated by the MPI Forum. Algorithms for executing collective operations, however, are not defined in the MPI standards. In a broadcast operation, all processes specify the same root process, whose buffer contents will be sent. Processes other than the root specify receive buffers. After the operation, all buffers contain the message from the root process.

A scatter operation, like the broadcast operation, is also a one-to-many collective operation. In a scatter operation, the logical root divides data on the root into segments and distributes a different segment to each compute node in the operational group (132). In scatter operation, all processes typically specify the same receive count. The send arguments are only significant to the root process, whose buffer actually contains sendcount*N elements of a given datatype, where N is the number of processes in the given group of compute nodes. The send buffer is divided and dispersed to all processes (including the process on the logical root). Each compute node is assigned a sequential identifier termed a 'rank.' After the operation, the root has sent sendcount data elements to each process in increasing rank order. Rank 0 receives the first sendcount data elements from the send buffer. Rank 1 receives the second sendcount data elements from the send buffer, and so on.

A gather operation is a many-to-one collective operation that is a complete reverse of the description of the scatter operation. That is, a gather is a many-to-one collective operation in which elements of a datatype are gathered from the ranked compute nodes into a receive buffer in a root node.

A reduction operation is also a many-to-one collective operation that includes an arithmetic or logical function performed on two data elements. All processes specify the same 'count' and the same arithmetic or logical function. After the reduction, all processes have sent count data elements from compute node send buffers to the root process. In a reduction operation, data elements from corresponding send buffer locations are combined pair-wise by arithmetic or logical operations to yield a single corresponding element in the root process' receive buffer. Application specific reduction operations can be defined at runtime. Parallel communications libraries may support predefined operations. MPI, for example, provides the following pre-defined reduction operations:

MPI_MAX	maximum
MPI_MIN	minimum
MPI_SUM	sum
MPI_PROD	product
MPI_LAND	logical and
MPI_BAND	bitwise and
MPI_LOR	logical or
MPI BOR	bitwise or
MPI_LXOR	logical exclusive or
MPI_BXOR	bitwise exclusive or

In addition to compute nodes, the parallel computer (100) includes input/output ('I/O') nodes (110, 114) coupled to compute nodes (102) through the high speed Ethernet network (174). The compute nodes (102) in the parallel com-

puter (100) may be partitioned into processing sets such that each compute node in a processing set is connected for data communications to the same I/O node. Each processing set, therefore, is composed of one I/O node and a subset of compute nodes (102). The ratio between the number of compute nodes to the number of I/O nodes in the entire system typically depends on the hardware configuration for the parallel computer (100). For example, in some configurations, each processing set may be composed of eight compute nodes and one I/O node. In some other configurations, each processing set may be composed of sixty-four compute nodes and one I/O node. Such example are for explanation only, however, and not for limitation. Each I/O node provides I/O services between compute nodes (102) of its processing set and a set of I/O devices. In the example of FIG. 1, the I/O nodes (110, 114) are connected for data communications I/O devices (118, 120, 122) through local area network ('LAN') (130) implemented using high-speed Ethernet.

The parallel computer (100) of FIG. 1 also includes a service node (116) coupled to the compute nodes through one of the networks (104). Service node (116) provides services common to pluralities of compute nodes, administering the configuration of compute nodes, loading programs into the compute nodes, starting program execution on the compute nodes, retrieving results of program operations on the compute nodes, and so on. Service node (116) runs a service application (124) and communicates with users (128) through a service application interface (126) that runs on computer terminal (122).

The parallel computer (100) of FIG. 1 operates generally for determining collective barrier operation skew in a parallel computer in accordance with embodiments of the present invention. A collective barrier operation is a function called by one node in an operational group, where the caller is blocked until all other nodes have also called the function. A barrier operation may be used to synchronize compute nodes that are otherwise executing independently and in parallel with one another. Barrier operation skew as the term is used here refers to an amount of time between a first completion of the collective barrier operation and the last completion of the collective barrier operation. In a non-homogenous system, compute nodes may take different amounts of time to enter and complete a collective barrier operation. The collective barrier skew is a measure of this difference.

The term 'non-homogenous' is used here to describe compute nodes having varying, or different, architectures or resources. Consider, for example, two non-homogenous compute nodes: a first compute node with 2 GB of DDR3 (Double Data Rate 3) RAM, a 1.8 GHz (Gigahertz) quad-core processor, and 1 gigabit per second (Gbps) Ethernet adapter; and a second compute node with 4 GB of DDR 2 (Double Data Rate 2) RAM, a 2.6 GHz dual-core processor, and a 10 Gbps Ethernet adapter. Each of these compute nodes when receiving a message, processing the message, and executing an computer program instruction responsive to the contents of the message, may complete the execution of the computer program instruction in different amounts of time. That is, even in instances in which the first and second compute node are sent an identical message at the same time, the resulting processing by each node may take a different amount of time. In this way, the compute nodes are said to be 'non-homogenous.' Homogenous nodes, by contrast, have the same or very similar architecture. Homogenous nodes are generally capable of receiving a message, processing the message, and executing computer program instructions in nearly the same amount of time or within some known, minor tolerance. In a

non-homogeneous system, barrier operation skew may be non-deterministic and may be exacerbated with respect to a homogenous system.

The parallel computer (100) of FIG. 1 operates generally for determining collective barrier operation skew in a parallel computer in accordance with embodiments of the present invention by carrying out an iterative measuring process—a process carried out once for each of the compute nodes until each compute node has been selected as a delayed node. The measuring process includes: selecting one of the plurality of compute nodes as a delayed node; entering, by each compute node other than the delayed node, a collective barrier operation; entering, after a predetermined delay by the delayed node, the collective barrier operation; receiving an exit signal from a root of the collective barrier operation; and measuring, for the delayed node, a barrier completion time. In the parallel computer (100) of FIG. 1, the barrier completion time is the period of time between the delayed node entering the barrier and receiving the exit signal. Once each compute node in the operational group has measured a barrier completion time, the barrier operation skew is calculated by: identifying, from the compute nodes' barrier completion times, a maximum barrier completion time and a minimum barrier completion time and calculating the barrier operation skew as the difference of the maximum barrier completion time and the minimum barrier completion time.

Determining collective barrier operation skew according to embodiments of the present invention is generally implemented on a parallel computer that includes a plurality of compute nodes organized for collective operations through at least one data communications network. In fact, such computers may include thousands of such compute nodes. Each compute node is in turn itself a kind of computer composed of one or more computer processing cores, its own computer memory, and its own input/output adapters. For further explanation, therefore, FIG. 2 sets forth a block diagram of an example compute node (102) useful in a parallel computer capable of determining collective barrier operation skew according to embodiments of the present invention. The compute node (102) of FIG. 2 includes a plurality of processing cores (165) as well as RAM (156). The processing cores (165) of FIG. 2 may be configured on one or more integrated circuit dies. Processing cores (165) are connected to RAM (156) through a high-speed memory bus (155) and through a bus adapter (194) and an extension bus (168) to other components of the compute node. Stored in RAM (156) is an application program (159), a module of computer program instructions that carries out parallel, user-level data processing using parallel algorithms.

Also stored in RAM (156) is a parallel communications library (161), a library of computer program instructions that carry out parallel communications among compute nodes, including point-to-point operations as well as collective operations. A library of parallel communications routines may be developed from scratch for use in systems according to embodiments of the present invention, using a traditional programming language such as the C programming language, and using traditional programming methods to write parallel communications routines that send and receive data among nodes on two independent data communications networks. Alternatively, existing prior art libraries may be improved to operate according to embodiments of the present invention. Examples of prior-art parallel communications libraries include the 'Message Passing Interface' ('MPI') library and the 'Parallel Virtual Machine' ('PVM') library.

Also stored in RAM (156) is an application (226). The application (226) in the example of FIG. 2 may be configured

as one instance of a parallel application with other instances executing amongst a plurality of compute nodes organized into an operational group. The application (226) in the example of FIG. 2 is configured for determining collective barrier operation skew in accordance with embodiments of the present invention by carrying out an iterative measuring process—a process carried out once for each of the compute nodes in the operational group until each compute node has been selected as a delayed node. During the iteration in which the compute node (102) is selected as the delayed node, the application (226) may take the following steps: entering, after a predetermined delay by the delayed node, a collective barrier operation—a barrier operation in which all other compute nodes in the operational group have entered. The application (226) may then receive an exit signal from a root of the collective barrier operation and measure, for the delayed node (102), a barrier completion time (228). The barrier completion time (228) is the period of time between the delayed node (102) entering the barrier and receiving the exit signal. Once each compute node in the operational group has measured a barrier completion time, the application (226) may then calculate a barrier operation skew (230) by: identifying, from the compute nodes' barrier completion times, a maximum barrier completion time and a minimum barrier completion time and calculating the barrier operation skew as the difference of the maximum barrier completion time and the minimum barrier completion time.

Also stored in RAM (156) is an operating system (162), a module of computer program instructions and routines for an application program's access to other resources of the compute node. It is typical for an application program and parallel communications library in a compute node of a parallel computer to run a single thread of execution with no user login and no security issues because the thread is entitled to complete access to all resources of the node. The quantity and complexity of tasks to be performed by an operating system on a compute node in a parallel computer therefore are smaller and less complex than those of an operating system on a serial computer with many threads running simultaneously. In addition, there is no video I/O on the compute node (102) of FIG. 2, another factor that decreases the demands on the operating system. The operating system (162) may therefore be quite lightweight by comparison with operating systems of general purpose computers, a pared down version as it were, or an operating system developed specifically for operations on a particular parallel computer. Operating systems that may usefully be improved, simplified, for use in a compute node include UNIX™, Linux™, Windows XP™, AIX™, IBM's i5/OS™, and others as will occur to those of skill in the art.

The example compute node (102) of FIG. 2 includes several communications adapters (172, 176, 180, 188) for implementing data communications with other nodes of a parallel computer. Such data communications may be carried out serially through RS-232 connections, through external buses such as USB, through data communications networks such as IP networks, and in other ways as will occur to those of skill in the art. Communications adapters implement the hardware level of data communications through which one computer sends data communications to another computer, directly or through a network. Examples of communications adapters useful in apparatus capable of determining collective barrier operation skew in a parallel computer include modems for wired communications, Ethernet (IEEE 802.3) adapters for wired network communications, and 802.11b adapters for wireless network communications.

The data communications adapters in the example of FIG. 2 include a Gigabit Ethernet adapter (172) that couples

example compute node (102) for data communications to a Gigabit Ethernet (174). Gigabit Ethernet is a network transmission standard, defined in the IEEE 802.3 standard, that provides a data rate of 1 billion bits per second (one gigabit). Gigabit Ethernet is a variant of Ethernet that operates over multimode fiber optic cable, single mode fiber optic cable, or unshielded twisted pair.

The data communications adapters in the example of FIG. 2 include a JTAG Slave circuit (176) that couples example compute node (102) for data communications to a JTAG Master circuit (178). JTAG is the usual name used for the IEEE 1149.1 standard entitled Standard Test Access Port and Boundary-Scan Architecture for test access ports used for testing printed circuit boards using boundary scan. JTAG is so widely adapted that, at this time, boundary scan is more or less synonymous with JTAG. JTAG is used not only for printed circuit boards, but also for conducting boundary scans of integrated circuits, and is also useful as a mechanism for debugging embedded systems, providing a convenient alternative access point into the system. The example compute node of FIG. 2 may be all three of these: It typically includes one or more integrated circuits installed on a printed circuit board and may be implemented as an embedded system having its own processing core, its own memory, and its own I/O capability. JTAG boundary scans through JTAG Slave (176) may efficiently configure processing core registers and memory in compute node (102) for use in dynamically reassigning a connected node to a block of compute nodes useful in systems capable of determining collective barrier operation skew in a parallel computer to embodiments of the present invention.

The data communications adapters in the example of FIG. 2 include a Point-To-Point Network Adapter (180) that couples example compute node (102) for data communications to a network (108) that is optimal for point-to-point message passing operations such as, for example, a network configured as a three-dimensional torus or mesh. The Point-To-Point Adapter (180) provides data communications in six directions on three communications axes, x, y, and z, through six bidirectional links: +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186).

The data communications adapters in the example of FIG. 2 include a Global Combining Network Adapter (188) that couples example compute node (102) for data communications to a global combining network (106) that is optimal for collective message passing operations such as, for example, a network configured as a binary tree. The Global Combining Network Adapter (188) provides data communications through three bidirectional links for each global combining network (106) that the Global Combining Network Adapter (188) supports. In the example of FIG. 2, the Global Combining Network Adapter (188) provides data communications through three bidirectional links for global combining network (106): two to children nodes (190) and one to a parent node (192).

The example compute node (102) includes multiple arithmetic logic units ('ALUs'). Each processing core (165) includes an ALU (166), and a separate ALU (170) is dedicated to the exclusive use of the Global Combining Network Adapter (188) for use in performing the arithmetic and logical functions of reduction operations, including an allreduce operation. Computer program instructions of a reduction routine in a parallel communications library (161) may latch an instruction for an arithmetic or logical function into an instruction register (169). When the arithmetic or logical function of a reduction operation is a 'sum' or a 'logical OR,' for example, the collective operations adapter (188) may

execute the arithmetic or logical operation by use of the ALU (166) in the processing core (165) or, typically much faster, by use of the dedicated ALU (170) using data provided by the nodes (190, 192) on the global combining network (106) and data provided by processing cores (165) on the compute node (102).

Often when performing arithmetic operations in the global combining network adapter (188), however, the global combining network adapter (188) only serves to combine data received from the children nodes (190) and pass the result up the network (106) to the parent node (192). Similarly, the global combining network adapter (188) may only serve to transmit data received from the parent node (192) and pass the data down the network (106) to the children nodes (190). That is, none of the processing cores (165) on the compute node (102) contribute data that alters the output of ALU (170), which is then passed up or down the global combining network (106). Because the ALU (170) typically does not output any data onto the network (106) until the ALU (170) receives input from one of the processing cores (165), a processing core (165) may inject the identity element into the dedicated ALU (170) for the particular arithmetic operation being performed in the ALU (170) in order to prevent alteration of the output of the ALU (170). Injecting the identity element into the ALU, however, often consumes numerous processing cycles. To further enhance performance in such cases, the example compute node (102) includes dedicated hardware (171) for injecting identity elements into the ALU (170) to reduce the amount of processing core resources required to prevent alteration of the ALU output. The dedicated hardware (171) injects an identity element that corresponds to the particular arithmetic operation performed by the ALU. For example, when the global combining network adapter (188) performs a bitwise OR on the data received from the children nodes (190), dedicated hardware (171) may inject zeros into the ALU (170) to improve performance throughout the global combining network (106).

For further explanation, FIG. 3A sets forth a block diagram of an example Point-To-Point Adapter (180) useful in systems capable of determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. The Point-To-Point Adapter (180) is designed for use in a data communications network optimized for point-to-point operations, a network that organizes compute nodes in a three-dimensional torus or mesh. The Point-To-Point Adapter (180) in the example of FIG. 3A provides data communication along an x-axis through four unidirectional data communications links, to and from the next node in the -x direction (182) and to and from the next node in the +x direction (181). The Point-To-Point Adapter (180) of FIG. 3A also provides data communication along a y-axis through four unidirectional data communications links, to and from the next node in the -y direction (184) and to and from the next node in the +y direction (183). The Point-To-Point Adapter (180) of FIG. 3A also provides data communication along a z-axis through four unidirectional data communications links, to and from the next node in the -z direction (186) and to and from the next node in the +z direction (185).

For further explanation, FIG. 3B sets forth a block diagram of an example Global Combining Network Adapter (188) useful in systems capable of determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. The Global Combining Network Adapter (188) is designed for use in a network optimized for collective operations, a network that organizes compute nodes of a parallel computer in a binary tree. The Global Combining Network Adapter (188) in the example of

11

FIG. 3B provides data communication to and from children nodes of a global combining network through four unidirectional data communications links (190), and also provides data communication to and from a parent node of the global combining network through two unidirectional data communications links (192).

For further explanation, FIG. 4 sets forth a line drawing illustrating an example data communications network (108) optimized for point-to-point operations useful in systems capable of determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. In the example of FIG. 4, dots represent compute nodes (102) of a parallel computer, and the dotted lines between the dots represent data communications links (103) between compute nodes. The data communications links are implemented with point-to-point data communications adapters similar to the one illustrated for example in FIG. 3A, with data communications links on three axis, x, y, and z, and to and fro in six directions +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186). The links and compute nodes are organized by this data communications network optimized for point-to-point operations into a three dimensional mesh (105). The mesh (105) has wrap-around links on each axis that connect the outermost compute nodes in the mesh (105) on opposite sides of the mesh (105). These wrap-around links form a torus (107). Each compute node in the torus has a location in the torus that is uniquely specified by a set of x, y, z coordinates. Readers will note that the wrap-around links in the y and z directions have been omitted for clarity, but are configured in a similar manner to the wrap-around link illustrated in the x direction. For clarity of explanation, the data communications network of FIG. 4 is illustrated with only 27 compute nodes, but readers will recognize that a data communications network optimized for point-to-point operations for use in determining collective barrier operation skew in a parallel computer in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes. For ease of explanation, the data communications network of FIG. 4 is illustrated with only three dimensions, but readers will recognize that a data communications network optimized for point-to-point operations for use in determining collective barrier operation skew in a parallel computer in accordance with embodiments of the present invention may in fact be implemented in two dimensions, four dimensions, five dimensions, and so on. Several supercomputers now use five dimensional mesh or torus networks, including, for example, IBM's Blue Gene Q™.

For further explanation, FIG. 5 sets forth a line drawing illustrating an example global combining network (106) useful in systems capable of determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. The example data communications network of FIG. 5 includes data communications links (103) connected to the compute nodes so as to organize the compute nodes as a tree. In the example of FIG. 5, dots represent compute nodes (102) of a parallel computer, and the dotted lines (103) between the dots represent data communications links between compute nodes. The data communications links are implemented with global combining network adapters similar to the one illustrated for example in FIG. 3B, with each node typically providing data communications to and from two children nodes and data communications to and from a parent node, with some exceptions. Nodes in the global combining network (106) may be characterized as a physical root node (202), branch nodes (204), and leaf nodes (206). The physical root (202) has two children but no parent

12

and is so called because the physical root node (202) is the node physically configured at the top of the binary tree. The leaf nodes (206) each has a parent, but leaf nodes have no children. The branch nodes (204) each has both a parent and two children. The links and compute nodes are thereby organized by this data communications network optimized for collective operations into a binary tree (106). For clarity of explanation, the data communications network of FIG. 5 is illustrated with only 31 compute nodes, but readers will recognize that a global combining network (106) optimized for collective operations for use in determining collective barrier operation skew in a parallel computer in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes.

In the example of FIG. 5, each node in the tree is assigned a unit identifier referred to as a 'rank' (250). The rank actually identifies a task or process that is executing a parallel operation according to embodiments of the present invention. Using the rank to identify a node assumes that only one such task is executing on each node. To the extent that more than one participating task executes on a single node, the rank identifies the task as such rather than the node. A rank uniquely identifies a task's location in the tree network for use in both point-to-point and collective operations in the tree network. The ranks in this example are assigned as integers beginning with 0 assigned to the root tasks or root node (202), 1 assigned to the first node in the second layer of the tree, 2 assigned to the second node in the second layer of the tree, 3 assigned to the first node in the third layer of the tree, 4 assigned to the second node in the third layer of the tree, and so on. For ease of illustration, only the ranks of the first three layers of the tree are shown here, but all compute nodes in the tree network are assigned a unique rank.

For further explanation, FIG. 6 sets forth a flow chart illustrating an example method determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. In the method of FIG. 6, the parallel computer includes a plurality of compute nodes organized into an operational group, such as a communicator in MPI. In some embodiments, the plurality of compute nodes in the operational group is non-homogenous. The term 'non-homogenous' is used here to describe compute nodes having varying, or different, architectures or resources. Consider, for example, two non-homogenous compute nodes: a first compute node with 2 GB of DDR3 (Double Data Rate 3) RAM, a 1.8 GHz (Gigahertz) quad-core processor, and 1 gigabit per second (Gbps) Ethernet adapter; and a second compute node with 4 GB of DDR 2 (Double Data Rate 2) RAM, a 2.6 GHz dual-core processor, and a 10 Gbps Ethernet adapter. Each of these compute nodes when receiving a message, processing the message, and executing a computer program instruction responsive to the contents of the message, may complete the execution of the computer program instruction in different amounts of time. That is, even in instances in which the first and second compute node are sent an identical message at the same time, the resulting processing by each node may take a different amount of time. In this way, the compute nodes are said to be 'non-homogenous.' Homogenous nodes, by contrast, have the same or very similar architecture. Homogenous nodes are generally capable of receiving a message, processing the message, and executing computer program instructions in nearly the same amount of time or within some known, minor tolerance.

A portion of the method of FIG. 6 is iterative—repeated once for each of the compute nodes in the operational group. In some embodiments the iterations may be carried out in order of the compute nodes' rank, beginning with the root of

13

the operational group. The iterative portion of the method of FIG. 6 includes selecting (602) one of the plurality compute nodes as a delayed node. Selecting (602) a compute node as a delayed node may be carried out in dependence upon compute node rank as described above. For example, in a first iteration, rank 0—the root—is selected as a delayed node; in a second iteration, rank 1 is selected as a delayed node; in the third iteration, rank 2 is selected as a delayed node; and so on until each compute node in the operational group has been selected as a delayed node.

The method of FIG. 6 also includes entering (604), by each compute node other than the delayed node, a collective barrier operation and entering (606), after a predetermined delay by the delayed node, the collective barrier operation. Effectively, all compute nodes other than the delayed node enter the barrier operation, while the delayed node waits. Then, after some delay that ensures all other compute nodes have entered the collective barrier operation, the delayed node enters the collective barrier. Upon entering the collective barrier operation, the delayed node may retrieve a current time as an entrance time to the collective barrier operation.

The method of FIG. 6 also includes receiving (608) an exit signal from a root of the collective barrier operation. An exit signal is a notification from a root compute node indicating all compute nodes have entered the global barrier operation. Each compute node in the collective barrier operation receives the exit signal and may then proceed executing other operations. In the method of FIG. 6, the delayed node may retrieve, upon receiving the exit signal, a current time as an exit time.

The method of FIG. 6 also includes measuring (610), for the delayed node, a barrier completion time, the barrier completion time. The barrier completion time in the example of FIG. 6 is a period of time between the delayed node entering the barrier and receiving the exit signal. The delayed node may measure (610) the barrier completion time by calculating the difference between the exit time and the entrance time. The barrier completion time is node-specific, meaning the time describes only one compute node's barrier operations.

The method of FIG. 6 continues by determining (612) whether each node has been selected as the delayed node. If each node has not been selected as the delayed node, the method of FIG. 6 repeats the steps of selecting (602) a delayed node; entering (604) a collective barrier operation by each compute node other than the delayed node; entering (606) the collective barrier operation by the delayed node after all other nodes; receiving (608) an exit signal, and measuring (610), for the delayed node, a barrier completion time.

Once each compute node has been selected as the delayed node—and thus each compute node has measured a node-specific barrier completion time, the method of FIG. 6 continues by calculating (614) the barrier operation skew. In the method of FIG. 6, calculating (614) the barrier operation skew includes identifying (616), from the compute nodes' barrier completion times, a maximum barrier completion time and a minimum barrier completion time and calculating (618) the barrier operation skew as the difference of the maximum barrier completion time and the minimum barrier completion time. A barrier operation skew as the term is used in this specification refers to an amount of time between a first completion of the collective barrier operation and the last completion of the collective barrier operation. In a non-homogenous system, compute nodes may take different amounts of time to enter and complete a collective barrier operation. The collective barrier skew is a measure of this difference.

14

The method of FIG. 6 also includes modifying (620) code that when executed carries out the collective barrier operation in dependence upon the barrier operation skew. That is, the collective barrier skew may be referenced during code development. A software developer may modify collective barrier operation code in order to minimize the collective barrier skew. The more minimized the skew, the more synchronized the compute nodes are upon completion of the barrier operation—a primary goal of the barrier operation itself.

For further explanation, FIG. 7 sets forth a line drawing illustrating an example operational group of compute nodes useful in determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. The operational group in the example of FIG. 7 includes seven compute nodes, each having a rank. One compute node having rank 0 is a root of the operational group. The operational group of FIG. 7 is configured as a binary tree structure comprising parent nodes and child nodes, with each parent node having no more than two children.

Determining collective barrier operation skew in a parallel computer that includes the operational group in FIG. 7 may be carried out by: selecting one of the plurality compute nodes as a delayed node; entering, by each compute node other than the delayed node, a collective barrier operation; entering, after a predetermined delay by the delayed node, the collective barrier operation; receiving an exit signal from a root of the collective barrier operation; and measuring, for the delayed node, a barrier completion time, the barrier completion time comprising the period of time between the delayed node entering the barrier and receiving the exit signal. These steps may be carried out iteratively, once for each compute node in the operational group.

Consider Table 1 below as an example of the result all iterations.

TABLE 1

Barrier Completion Time Measurements For The Operational Group Of FIG. 7	
Delayed Node Rank ID	Barrier Completion Time
0	2 ms
1	7 ms
2	10 ms
3	8 ms
4	3 ms
5	12 ms
6	9 ms
7	5 ms

Each record in the example Table 1 above relates to an iteration of selecting a delayed node; entering, by each compute node other than the delayed node, a collective barrier operation; entering the collective barrier operation by the delayed node after all other nodes; receiving an exit signal; and measuring, for the delayed node, a barrier completion time. In a first iteration, rank 0 measured a barrier completion time of 2 ms, in a second iteration, rank 1 measured a barrier completion time of 7 ms and so on.

Once each compute node has been selected as the delayed node and measured a node-specific barrier completion time, the barrier operation skew may be calculated by identifying, from the compute nodes' barrier completion times, a maximum barrier completion time and a minimum barrier completion time and calculating the barrier operation skew as the difference of the maximum barrier completion time and the minimum barrier completion time. Continuing with the barrier completion times in the example of Table 1 above, the

15

maximum barrier completion time is 12 ms, measured by the compute node having rank 5, and the minimum barrier completion time is 2 ms, measured by the compute node having rank 0. The barrier operation skew in this example, is 10 ms, the difference between the maximum (12 ms) and the minimum (2 ms) completion times.

For further explanation, FIG. 8 sets forth a flow chart illustrating a further example method determining collective barrier operation skew in a parallel computer according to embodiments of the present invention. The method of FIG. 8 is similar to the method of FIG. 6 in that the parallel computer within which the method of FIG. 8 is carried out also includes plurality of compute nodes organized into an operational group, such as a communicator in MPI. The method of FIG. 8 is also similar to the method of FIG. 6 in that the method of FIG. 8 includes: iteratively, for each compute node: selecting (602) a delayed node; entering (604), by each compute node other than the delayed node, a collective barrier operation; entering (606), after a predetermined delay by the delayed node, the collective barrier operation; receiving (608) an exit signal; measuring (610), for the delayed node, a barrier completion time; and calculating (614) the barrier operation skew.

The method of FIG. 8 differs from the method of FIG. 6, however, in that in the method of FIG. 8 entering (604), by each compute node other than the delayed node, a collective barrier operation is carried out by executing (802), by the delayed node, a predefined delay operation. The execution of the predefined delayed operation is configured to be carried out in a period of time greater than a period of time required for other nodes to enter the barrier operation. That is the delayed node executes an operation configured such that during its execution all other nodes have enough time to enter the barrier.

Also in the method of FIG. 8, measuring (610) a barrier completion time may include storing (804) the barrier completion time locally at the delayed node. That is, each node, upon measuring its node-specific barrier completion time, stores that value locally until all nodes have computed a node-specific barrier completion time. Then, when identifying (616) a maximum and minimum barrier completion time from the node's barrier completion times, the root node may execute, among the compute nodes, one or more ALLREDUCE operations. The root node, for example may initiate an ALLREDUCE operation that specifies, as an operation to perform during the ALLREDUCE, a 'max' operation followed by an ALLREDUCE operation that specifies as an operation to perform during the ALLREDUCE, a 'min' operation. In this way, the maximum barrier time and minimum barrier time may be calculating in one or more collective, parallel operations.

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable transmission medium or a computer readable storage medium. A computer readable storage

16

medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

A computer readable transmission medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable transmission medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to

17

function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

It will be understood from the foregoing description that modifications and changes may be made in various embodiments of the present invention without departing from its true spirit. The descriptions in this specification are for purposes of illustration only and are not to be construed in a limiting sense. The scope of the present invention is limited only by the language of the following claims.

What is claimed is:

1. A method of determining collective barrier operation skew in a parallel computer comprising a plurality of compute nodes organized into an operational group, the method comprising:

for each of the compute nodes until each compute node has been selected as a delayed node:

selecting one of the plurality compute nodes as a delayed node;

entering, by each compute node other than the delayed node, a collective barrier operation including sending to a root of the collective barrier operation, by each compute node, an indication that the compute node has finished the collective barrier operation and is waiting for an exit signal;

18

entering, after a predetermined delay by the delayed node, the collective barrier operation including sending to the root of the collective barrier operation, by the delayed node, an indication that the delayed node has finished the collective barrier operation and is waiting for the exit signal;

receiving from each compute node, by the root, an indication that the compute node has finished the collective barrier operation and is waiting for the exit signal, wherein the indication from the delayed node is received last after receiving the indications from all of the other compute nodes;

in response to receiving all of the indications from all of the compute nodes, transmitting, by the root node, the exit signal to each of the compute nodes;

measuring, for the delayed node, a barrier completion time, the barrier completion time comprising the period of time between the delayed node entering the collective barrier operation and receiving the exit signal; and

calculating the barrier operation skew including: identifying, from the compute nodes' barrier completion times, a maximum barrier completion time and a minimum barrier completion time; and calculating the barrier operation skew as the difference of the maximum barrier completion time and the minimum barrier completion time; and

modifying code that when executed carries out the collective barrier operation in dependence upon the calculated barrier operation skew.

2. The method of claim 1, wherein the plurality of compute nodes are non-homogenous.

3. The method of claim 1, wherein:

entering, after a predetermined delay by the delayed node, the collective barrier operation further comprises: retrieving, upon entering the collective barrier operation, a current time as an entrance time;

receiving an exit signal from a root of the collective barrier operation further comprises retrieving, by the delayed node upon receiving the exit signal, a current time as an exit time; and

measuring a barrier completion time further comprises calculating, as the barrier completion time, the difference between the exit time and the entrance time.

4. The method of claim 1, wherein:

measuring a barrier completion time further comprises storing the barrier completion time locally at the delayed node; and

identifying, from the node's barrier completion times, a maximum barrier completion time and a minimum barrier completion time further comprises executing, among the compute nodes, an ALLREDUCE operation.

* * * * *